

VOLUME AND GEOMETRY OF HOMOGENEOUSLY ADEQUATE KNOTS

PAIGE BARTHOLOMEW, SHANE MCQUARRIE, JESSICA S. PURCELL, AND KAI WESER

ABSTRACT. We bound the hyperbolic volumes of a large class of knots and links, called homogeneously adequate knots and links, in terms of their diagrams. To do so, we use the decomposition of these links into ideal polyhedra, developed by Futer, Kalfagianni, and Purcell. We identify essential product disks in these polyhedra.

1. INTRODUCTION

A knot or link in S^3 is often represented by a diagram: a planar 4-valent graph with over-under crossing information. Results from hyperbolic geometry tell us that the link complement is often hyperbolic [19], and invariants of the hyperbolic geometry, such as the volume, are knot and link invariants [14, 16]. However, it is a nontrivial problem to bound the volume of a hyperbolic knot in terms of its diagram. To date, this has been done for many classes of knots. Lackenby found conditions that bound the volume of alternating links based on their diagrams [11]. Purcell [17], and Futer, Kalfagianni, and Purcell found volume bounds of many additional classes of links [5, 6, 7]. Most recently, they found bounds on the volumes of a class of links called semi-adequate links (defined in Definition 2.3) in terms of their diagrams [3]. These links include alternating links, and many additional classes of links.

Much of the work in [3] was extended to an even larger class of links, called homogeneously adequate links. However, the step of bounding volumes in terms of diagrams was not completed in that paper. It was shown that volumes could be bounded below by the Euler characteristic of a graph coming directly from the diagram, but that a certain quantity must be subtracted from this result. This quantity, called $||E_c||$ in that paper, was shown to have relations to diagrams for semi-adequate links, but the connection was not generalized to homogeneously adequate links.

In this paper, we investigate this quantity for homogeneously adequate links. We prove that it is bounded in terms of properties of the diagram. Thus we can extend many results on volumes from [3] to homogeneously adequate links.

Why homogeneously adequate links? These links, defined carefully below, are known to include nearly all classes of knots and links mentioned above. In particular,




FIGURE 1. From left to right: an unresolved crossing, the A -resolution, the B -resolution.

they include alternating links, positive and negative closed braids, closed 3-braids, Montesinos links, and homogeneous links. They include links that are known not to be semi-adequate, including the closed 5-braid shown in Figure 3 in this paper. See [9, Remark 1.5]. At the time of this writing, it is unknown whether *every* hyperbolic link admits a homogeneously adequate diagram. See [15] for a detailed discussion.

1.1. Organization. In Section 2, we review the definition of homogeneously adequate links, and the polyhedral decomposition obtained for these links in [3]. In Section 3, we review some of the results of [3] that we will need in this paper, including a few of the arguments using a combinatorial proof technique that we call “tentacle chasing”. We prove the main technical result in Section 4. This leads to bounds on the quantity $||E_c||$ in terms of properties of the diagram. The results of that section are applied to volumes in Section 5.

1.2. Acknowledgements. This work was partially supported by NSF grant DMS-1252687, and by funding for student research from Brigham Young University’s College of Physical and Mathematical Sciences.

2. PROPERTIES OF KNOTS

In this section, we review the definitions of semi-adequate and homogeneously adequate knots and links. We also describe the polyhedral decomposition developed in Chapters 2 and 3 of [3]. In order to do so, we need to review a few other definitions. Thus this section is particularly heavy with definitions and notation. We attempt to include enough examples to make this paper as self-contained and readable as possible. However, in the interest of space and readability, we must omit many of the details in [3], including the geometric ideas that led to the development of these polyhedra. For more information on these ideas, and more information on the polyhedra and their other uses, see, for example, [8], in addition to [3].

Given a crossing of a knot, there are two ways to resolve the crossing, called the A -resolution and the B -resolution. See Figure 1.

Definition 2.1. A choice of A - or B -resolution for each crossing of the diagram is a *state*, which we will denote by σ . After applying σ to the diagram, we obtain a crossing-free diagram consisting of a collection of closed curves called *state circles*, which we denote by s_σ . In Figure 1, there is a dashed edge running between state circles in each resolution. We will add these edges to the state circles in s_σ , obtaining




FIGURE 2. Left to right: A diagram, the graph H_A (i.e. the state σ is the all- A state), the graph \mathbb{G}_A , the graph \mathbb{G}'_A , and the graph H_B (i.e. the state is the all- B state)

a 3-valent graph, which we denote by H_σ . The edges added to go from s_σ to H_σ we will call *segments*.

Definition 2.2. The *state graph* \mathbb{G}_σ is obtained from H_σ by collapsing each state circle to a vertex. Hence edges of \mathbb{G}_σ are segments of H_σ . The *reduced state graph* \mathbb{G}'_σ is obtained from \mathbb{G}_σ by removing all multiple edges between pairs of vertices.

For an example of a diagram and graphs H_σ , \mathbb{G}_σ , and \mathbb{G}'_σ for the case σ is the all- A state, see Figure 2, part of which appeared in [8]. For an example of a diagram and H_σ for a more general state, see Figure 3.

Definition 2.3. A diagram $D(K)$ is called σ -adequate if \mathbb{G}_σ has no one-edge loops. That is, there is no edge of \mathbb{G}_σ connecting a vertex to itself.

A knot or link is called *semi-adequate* if it admits a diagram for which the all- A or all- B state σ is adequate.

Note that the link in Figure 2 is A -adequate, but not B -adequate, because a segment at the top of the graph of H_B connects a state circle to itself.

Semi-adequate links appeared first in the study of Jones-type invariants [12, 18], and consist of a broad class of links. These links were the main focus of study in [3]. However, that paper also studied broader classes of links as below, which are the focus of this paper.

Definition 2.4. The state circles s_σ cut the plane into regions, and segments are placed in these regions to form H_σ . We say that a state σ is *homogeneous* if for any region in the complement of s_σ , all segments in that region in H_σ come from A -resolutions, or all segments in that region come from B -resolutions.

A knot or link is *homogeneous* if it admits a diagram and a state σ that is homogeneous.

Note that the all- A and all- B states will always be homogeneous, since they are composed exclusively of A -resolutions or B -resolutions, respectively. Figure 3




FIGURE 3. A link (represented as a closed 5-braid) and H_σ for a homogeneous state σ .

shows a diagram of a link and a state σ that is homogeneous, but more general than the all- A or all- B state.

Definition 2.5. A *homogeneously adequate link* is a link that admits a diagram and a choice of state σ that is both adequate and homogeneous.

In the examples above, the all- A state of Figure 2 is both adequate and homogeneous, and the state σ shown in Figure 3 is also adequate and homogeneous. Hence these are both homogeneously adequate links.

In [3], the following theorem was obtained bounding volumes of homogeneously adequate links.

Theorem 2.6 (Theorem 9.3 of [3]). *Let K be a link with a prime diagram that admits a homogeneously adequate state σ . Then the volume of $S^3 - K$ satisfies*

$$\text{vol}(S^3 - K) \geq v_8(\chi_-(\mathbb{G}'_\sigma) - \|E_c\|),$$

where $v_8 = 3.6638\dots$ is the volume of a hyperbolic regular ideal octahedron, and $\chi_-(\cdot)$ denotes the negative Euler characteristic, or $\max\{-\chi(\cdot), 0\}$. Finally, $\|E_c\|$ counts complex essential product disks in a spanning set.

The term $\|E_c\|$ will be defined carefully below. For now, note that the term $\chi_-(\mathbb{G}_\sigma)$ can easily be read off the prime, homogeneously adequate diagram of the link K . It is more difficult to determine $\|E_c\|$. Nevertheless, we show that $\|E_c\|$ is bounded by properties of the diagram of a homogeneously adequate link. To do so, we will need to review the polyhedral decomposition in [3].

2.1. Polyhedral Decomposition. In [3], it was shown how to turn a homogeneously adequate diagram of a link into a collection of checkerboard ideal polyhedra. Not all the details of the decomposition into polyhedra will be necessary here. However, we will review properties of the polyhedra that we use later in this paper.

Given a state σ , recall that the state circles s_σ cut the plane of projection into distinct regions. We say a region is *nontrivial* if it contains more than one state




FIGURE 4. Left to right: A graph H_σ ; a nontrivial region; the associated alternating sub-diagram (assuming the region is an A -region); a maximal collection of non-prime arcs. The shaded area indicates one polyhedral region (of four polyhedral regions total).

circle on its boundary. In Figure 4, an example of a graph H_σ is shown on the left, as well as a nontrivial region, second from the left.

Let $D(K)$ be a diagram and σ a state such that σ is homogeneously adequate. For any nontrivial region in the complement of s_σ , one may form an alternating sub-diagram of $D(K)$ by starting with the state circles of s_σ on the boundary of that region, and adding to the state circles the crossings associated with segments contained inside that region. Figure 4, third from left, shows an example. Note that homogeneity of the diagram ensures that the result is alternating. However, the resulting alternating subdiagram may not be prime.

Definition 2.7. Recall that a diagram is *prime* if every simple closed curve meeting the diagram transversely in exactly two edges bounds a region of the diagram with no crossings.

Menasco described a polyhedral decomposition for prime alternating diagrams [13], which we wish to use. If our alternating subdiagram is not prime, there is an arc in H_σ , running through the associated nontrivial region, with both endpoints on the same state circle C , which arc meets no other state circles, but splits the given region into two sub-regions, both containing state circles besides C .

Definition 2.8. An arc as above is called a *non-prime arc*. That is, it is an arc embedded in the complement of H_σ with both endpoints on a state circle C , splitting a region of the complement of s_σ into two sub-regions, both of which contain segments of H_σ .

We say a collection of non-prime arcs $\alpha_1, \dots, \alpha_n$ is *maximal* if, for each nontrivial component in the complement of s_σ , the union of the arcs $\alpha_1, \dots, \alpha_n$ split the corresponding component of H_σ into components, each of which contains segments, and moreover, the addition of any other non-prime arc yields such a component that does not contain segments.

A maximal collection of non-prime arcs is shown in the above example, Figure 4, right.

Definition 2.9. Let $\alpha_1, \dots, \alpha_n$ be a maximal collection of non-prime arcs. A region in the complement of the union of the state circles s_σ and the non-prime arcs $\bigcup_i \alpha_i$ is called a *polyhedral region*.

See Figure 4 for an example. As in this example, typically we think of a polyhedral region as living in H_σ . That is, add to the region the segments of H_σ . Note that because σ is a homogeneous state, all the segments will come from A -resolutions, or all of them will come from B -resolutions in the polyhedral region.

Definition 2.10. Given a polyhedral region, form an associated prime alternating diagram by taking circles corresponding to the boundary of the region and adding to these a crossing for each segment in the region. Associated to any such alternating diagram is a checkerboard colored ideal polyhedron, obtained by taking the polyhedron with the diagram graph on its boundary, and vertices removed (at infinity). This is the usual checkerboard colored ideal polyhedron associated with a prime alternating link, as described by Menasco [13]. We call it the *lower polyhedron* associated with the polyhedral region.

The lower polyhedra form the portions of our polyhedral decomposition that are easiest to understand, since they are standard polyhedra of alternating links. They don't play much of a role in this paper. There is one additional, more complicated polyhedron, the *upper polyhedron*, which plays the most important role in this work. Although it is more complicated than the lower polyhedra, all the properties of the upper polyhedron can be read off of the graph H_σ . We spend some time now describing the upper polyhedron in terms of H_σ .

The upper polyhedron consists of white faces and shaded faces. White faces will be the components in the complement of shaded faces.

The shaded faces are made up of three pieces, namely *innermost disks*, *tentacles*, and *non-prime switches*. We first define these three parts, then describe how they fit together.

Definition 2.11. Each trivial region in the complement of s_σ is bounded by a single state circle. We call a trivial region an *innermost disk*.

Definition 2.12. A *tentacle* is a portion of shaded face running adjacent to exactly one segment and state circle. A tentacle is contained in a single polyhedral region. If the polyhedral region is an A -region, meaning all segments come from A -resolutions, then the tentacles in that region have the shape of an L : The *head*, or vertical portion of the L , lies against the segment, and the *tail*, or horizontal portion of the L , lies against the state circle. The tail terminates when it meets a segment lying on the same side of the state circle. We say a tentacle in an A -region runs down and right, or right-down. Figure 5, middle, shows right-down tentacles in an example.

If the polyhedral region is a B -region, then the tentacles in the region have the shape of a Γ : The *head*, or vertical portion of the Γ , lies against a segment. The *tail*, or horizontal portion, runs adjacent to a state circle. It terminates when it meets another segment on the same side. We say a tentacle in a B -region runs down and left, or left-down.

Definition 2.13. For any non-prime arc, a *non-prime switch* is a portion of a shaded face forming a regular neighborhood of the non-prime arc in H_σ .

Definition 2.14. Starting with the graph of H_σ , form the *upper polyhedron* as follows.

- (1) Remove bits of state circles at the top-right and bottom-left of A -segments and at the top-left and bottom-right of B -segments, as in Figure 5, left.
- (2) Add maximal collection of non-prime arcs $\alpha_1, \dots, \alpha_n$ (shown as dashed red lines in Figure 5 left).
- (3) Color each innermost disk a unique color.
- (4) Add tentacles as follows: Tentacles exit state circles at gaps and flow right-down into A -regions or left-down into B -regions, and they are given the same color as the shaded face on the opposite side of the initial gap. They terminate when they meet another segment. See Figure 5, middle.
- (5) Join tentacles by non-prime switches, merging two different colored faces into one, as in Figure 5, right.

The above process defines the shaded faces of the upper polyhedron. In [3], it was shown that the shaded faces are simply connected.

The rest of the upper polyhedron consists of the following data.

- The white faces are the remaining, unshaded regions of the projection plane. They are in one-to-one correspondence with regions in the complement of $H_\sigma \cup (\cup_i \alpha_i)$.
- Ideal edges separate white and shaded faces.
- Ideal vertices are remnants of H_σ . Note they zig-zag in H_σ , following portions of state circles and segments. We often call them *zig-zag vertices*.

In [3], it was shown that the upper and lower polyhedra give a decomposition of the manifold, denoted M_σ , obtained by cutting the knot complement along a particular incompressible surface, the *state surface* S_σ . For this paper, we do not need a description of that surface or the cut manifold M_σ , only the results that follow from the polyhedral decomposition, and so we omit the details here.

A shaded face of the upper polyhedron determines a directed graph. Innermost disks are sources, which correspond to vertices of the directed graph. Tentacles can be represented by directed edges running from head to tail. A non-prime switch is given by four directed edges, running in the direction of adjacent tentacles, and a vertex. Starting in Section 3, we will frequently consider directed arcs inside of




FIGURE 5. Constructing shaded faces. Left: From H_σ , remove bits of state circles (solid lines) and add in non-prime arcs (dashed). Middle: Color the innermost disks and extend tentacles. Right: Add non-prime switches.

shaded faces, and think of them as running through the directed graph. We need a couple of definitions to explain their direction. We also want to ensure the arcs we consider are simplified, in the sense that they don't double back on themselves. These are the reasons for the following definitions.

Definition 2.15. We say that an arc running monotonically through the directed graph of a shaded face is *simple with respect to the shaded face*. Any arc in a shaded face may be isotoped to be simple.

A simple arc running through a tentacle is said to run *downstream* if it runs in the direction of the tentacle, i.e. from head to tail. Otherwise, it runs *upstream*. A simple arc running through a shaded face can only change direction, downstream to upstream, inside a non-prime switch.

We are now ready to define the complex essential product disks of Theorem 2.6. We only need to describe the notion of an essential product disk in the upper polyhedron.

Definition 2.16. In the upper polyhedron, an *essential product disk (EPD)* is a properly embedded essential disk whose boundary meets distinct ideal vertices of the upper polyhedron (remnants of H_σ) exactly twice, and which lies completely in two colored faces otherwise.

Definition 2.17. An EPD in the upper polyhedron is said to be *simple* if its boundary encloses a single white bigon face of the polyhedron. It is said to be *semi-simple* if its boundary encloses a union of white bigon faces to one side, and no other white faces to that side. It is *complex* otherwise.

Definition 2.18. A finite collection of disjoint EPDs E_1, \dots, E_n *spans* a set B if $B \setminus (E_1 \cup \dots \cup E_n)$ is a finite collection of prisms (I -bundles over a polygon) and solid tori.

It was shown in [3] that a characteristic I -bundle associated to the knot complement is spanned by EPDs completely contained in the ideal polyhedra of the polyhedral decomposition. Moreover, there exists a set $E_s \cup E_c$ of EPDs in the upper polyhedron P_σ that spans the intersection of the I -bundle with P_σ , [3, Lemma 5.8]. Here E_s consists of all simple disks in the upper polyhedron, and E_c consists of complex disks. The collection E_c is minimal in the following sense.

Suppose the boundary of an EPD runs adjacent to both sides of an ideal vertex without meeting that vertex. Then, by connecting the boundary across this vertex, and surgering the EPD, we obtain two new EPDs. We say the two EPDs are obtained by a *parabolic compression* of the original. The collection E_c is minimal in the sense that no EPD in E_c parabolically compresses to a subcollection of $E_s \cup E_c$.

Definition 2.19. The term $||E_c||$ in Theorem 2.6 is the number of complex EPDs in the collection E_c .

We wish to bound $||E_c||$ in terms of the diagram.

3. USEFUL LEMMAS FOR TENTACLE CHASING

Our goal now is to find EPDs in the upper polyhedron. In order to do so, we will use combinatorics of the graph H_σ and combinatorics of the upper polyhedron, which we visualize as lying on the graph H_σ . Many of our tools come from [3]. In this section, we review the main tools we will be using.

The first set of results come from [3, Chapter 3]. In that paper, these lemmas are called Escher Stairs, Shortcut Lemma, Downstream Lemma, and the Utility Lemma. The proofs in that paper were originally written in the semi-adequate case rather than the homogeneously adequate case. However, it was observed there that the proofs carried over into the homogeneously adequate case (see section 3.4 of that paper). We state the results here for reference, and include brief proofs as a warm up to the combinatorial proof technique of tentacle chasing. First, we need a definition.

Definition 3.1. A *staircase* is a connected subgraph of H_σ consisting of finitely many segments alternating with portions of state circles, which has the following properties. First, two segments with a state circle between them lie on opposite sides of the state circle. Second, segments and state circles in the staircase are given directions, such that the graph runs monotonically from “top” to “bottom.” A *step* consists of a segment and the portion of state circle in the direction of the segment. If the segment lies in an A -region, then the state circle runs right from the end of the step. If it lies in a B -region, it runs left. See Figure 6.

Lemma 3.2 (Escher stairs). *Let K be a link with a diagram and a state σ that is homogeneous and adequate. In the graph H_σ , the following hold.*




FIGURE 6. Left to right: A staircase with all segments in A -regions. A staircase with all segments in B -regions. A mixed staircase.




FIGURE 7. Left: a right-down staircase forming a loop. Right: a right-down staircase with top and bottom on same side of the same state circle.

- (1) *There is no staircase in which one state circle C is both the top and the bottom of the staircase, with top and bottom steps on the same side of C .*
- (2) *No staircase forms a loop, i.e. top and bottom steps lie on opposite sides of the same state circle.*

See Figure 7.

Proof. Suppose by way of contradiction that such a staircase exists. Notice that the segments and state circles of the staircase form a closed curve γ in the diagram. The state circles of the staircase intersect γ . Since the state circles are also closed curves, they must connect to other state circles of the staircase on both sides of γ . We show by induction on the number of stairs of the staircase that this situation inevitably violates the adequacy of the diagram.

Consider the first result: the staircase in which one state circle C is both the top and bottom, with top and bottom segments on the same side of C . Let C_1, \dots, C_n denote the other state circles in the staircase. If n is odd, then not all state circles of the staircase can connect to each other, violating our observation above. So n must be even. If $n = 2$, C_1 and C_2 must be connected. However, there is a segment of the staircase between C_1 and C_2 , contradicting the adequacy of this state. See Figure 8. Therefore, there cannot exist a staircase with only three state circles in which one of the state circles is the top and bottom of the staircase.

We now proceed with the inductive step. Assume by induction that it is impossible to form a staircase with top and bottom on the same step with p state circles, for $p = 2, 3, \dots, n$. We show that for $p = n + 1$ the scenario is also impossible. Let the state circles be denoted $C, C_1, C_2, \dots, C_{n-1}, C_n$, with C at the top and bottom. As before, we must connect the loose ends of the state circles C_i enclosed by γ . The




FIGURE 8. Examples of staircases with three state circles, and top and bottom on the same state circle.




FIGURE 9. Left: a staircase with $n + 1$ state circles with C at top and bottom. Right: there exists a sub-staircase with C_n at the top and bottom.

state circle C_n must connect to some state circle C_j inside γ . This connection cuts the staircase into sub-staircases, one of which is a staircase in which the state circle $C_n = C_j$ is the top and bottom of the staircase. See Figure 9. Now the number of state circles in this sub-staircase is at most $n - 2$. This is impossible.

We now prove the second result of the lemma, that no staircase may form a loop. Suppose that C_1, \dots, C_n are state circles in a staircase forming a loop, with top and bottom segments on opposite sides of C_1 . The staircase forms a simple closed curve γ in the plane. Hence state circles C_1, \dots, C_n must connect to each other on either side of γ . Then C_n must connect to some state circle C_j . But now the sub-staircase from C_j to C_n has top and bottom segments on the same state circle $C_n = C_j$, on the same side. This contradicts our previous result. \square

Recall that a non-prime arc has endpoints on a state circle C , and divides a region bounded by C in the plane into two regions. We say that each such region is a *non-prime half-disk*.

Lemma 3.3 (Shortcut lemma). *Let α be a non-prime arc with endpoints on a state circle C in a diagram with homogeneous, adequate state σ . Suppose a directed arc γ lies entirely inside a single shaded face, and is simple with respect to that shaded face. Suppose γ runs across α into the interior of the non-prime half-disk bounded by α and C , and then runs upstream. Finally, suppose that γ exits the interior of that half-disk across the state circle C . Then γ must exit by following a tentacle downstream (that is, it cannot exit running upstream).*

Proof. Consider an innermost counterexample. That is, assume that γ exits the non-prime half-disk bounded by non-prime arc α by running upstream, but if γ enters

into any other non-prime half-disk after crossing α and before exiting the half-disk bounded by α , then it exits running downstream. According to the hypothesis, upon crossing non-prime arc α , γ turns upstream through a tentacle. Note that this tentacle may be in either an A or B region. Once γ has begun running upstream, γ may continue travelling in the following ways: (a) upstream, or (b) downstream, including the case that it runs into a non-prime half-disk.

For the first case, γ continues upstream until it meets another state circle and encounters the same options of continuation. By Lemma 3.2, γ may not exit α by continuing to run upstream. Thus γ must, at some point, run downstream.

Consider the second case. We claim that once γ begins running downstream it must continue downstream until it exits the state circle C . If γ crosses a state circle going downstream, it will still be running downstream in the new tentacle, because it is simple. If γ crosses into another non-prime half-disk it must eventually exit running downstream, even if it runs upstream at some point within the half-disk, by the assumption that α is the innermost counterexample. The arc γ may also run over a non-prime arc without entering the half-disk it bounds, but this results in no change of direction.

Therefore, in all cases, γ must exit α by crossing C by running downstream. \square

Lemma 3.3 implies the following, stronger lemma, whose proof we omit, but is contained in [3].

Lemma 3.4 (Downstream). *Let γ be a directed arc lying in a single shaded face that is simple with respect to the shaded face. Suppose γ begins by crossing a state circle running downstream, and suppose that every time γ crosses a non-prime arc, it exits the corresponding non-prime half-disk. Then γ defines a staircase such that every segment of the staircase is adjacent to γ , with γ running downstream. Moreover, the endpoints of γ lie on tentacles adjacent to the first and last stairs.*

Finally, the following lemma is most useful for tentacle chasing arguments.

Lemma 3.5 (Utility lemma). *Let γ be a simple directed arc in a shaded face that starts and ends on the same state circle C . Then the following are true:*

- (1) γ starts by running upstream away from C ,
- (2) γ terminates running downstream toward C , and
- (3) γ cannot intersect C more than two times.

Proof. We will consider three results individually and prove each by contradiction.

For item (1), assume that γ begins by running downstream, contrary to the claim. By Lemma 3.4, there exists a staircase starting on C such that γ runs downstream, adjacent to each segment of the staircase. Since γ starts and ends on C , the staircase must start and end on C . This contradicts Lemma 3.2 (Escher Stairs).

For item (2), we know γ begins at C running upstream. Assume that γ terminates at C running upstream. By reversing the orientation of γ , we obtain the same scenario as in part (1). Therefore, γ must terminate running downstream.

For item (3), assume γ crosses C more than once, at points x_1, x_2, \dots, x_n . Consider the sub arc of γ from x_1 to x_2 . By item (2), the subarc must cross C going downstream at x_2 . However, this means that for the subarc of γ from x_2 to x_3 , the arc travels downstream from C at x_2 , which contradicts item (1). Therefore, γ may only cross C two times. \square

4. ESSENTIAL PRODUCT DISKS

The following theorem, whose proof is contained in this section, is the main result of this paper. It generalizes [3, Theorem 6.4], which was proved only for A -adequate links. Our proof is similar to the proof in [3], following the same sequence of steps. However we need to include more cases when we allow both A and B resolutions in homogeneously adequate diagrams.

Theorem 4.1. *Let $D(K)$ be a prime, σ -adequate σ -homogeneous diagram of a link K in S^3 . Denote by P_σ the upper polyhedron in the prime polyhedral decomposition described above. Suppose there is an essential product disk E embedded in P_σ . Then there is a 2-edge loop in \mathbb{G}_σ so that ∂E runs over tentacles adjacent to segments of the 2-edge loop.*

The A -adequate versions of many of the following lemmas appear in [3, Chapter 6]. We begin with a lemma which allows us to work with normal squares in the upper polyhedron, rather than EPDs. This is the analogue to [3, Lemma 6.1], and its proof is identical.

Lemma 4.2 (EPD to oriented square). *Let $D(K)$ be a prime, σ -adequate and σ -homogeneous diagram of a link in S^3 . Let P_σ denote the upper polyhedron. Suppose there is an EPD E properly embedded in P_σ . Then ∂E can be pulled off the ideal vertices of P_σ to give a normal square E' in P_σ with the following properties.*

- (1) *Two opposite edges of $\partial E'$ run through shaded faces, which we label green and orange.*
- (2) *The other two opposite edges run through white faces, each cutting off a single vertex of the white face.*
- (3) *The single vertex cut off in a white face forms a triangle in that white face, oriented in such a way that in counter-clockwise order, the edges of the triangle are colored orange–green–white.*

Finally, the two white edges of the normal square cannot lie on the same white face of the polyhedron.

Proof. The EPD E runs over two shaded faces, labeled green and orange, and runs between them over two ideal vertices. If we pull ∂E off of the ideal vertex slightly,




FIGURE 10. Type \mathcal{A} : EPD runs over tentacles of different colors adjacent to 2-edge loop, white side cuts off a triangle with vertex (zig-zag) containing one of the segments of the loop.

into one of the two white faces, the result will cut off a triangle. The orientation is determined by which side of the vertex ∂E is pulled into, and we choose the orientation that matches item (3) of the lemma. The white faces must be distinct else the boundary of the normal square E' on one of the shaded faces can be joined to an arc between the two white sides of $\partial E'$ to form a normal bigon, contradicting the fact that P_σ admits no normal bigons [3, Proposition 3.24]. \square

In the above lemma, the term *normal* is used without definition. A precise definition is given in [3, Definition 3.14]. However, we will only work with normal squares that have the form of Lemma 4.2.

Remark 4.3. When a normal square has the form of Lemma 4.2, note that the white faces will lie at the tails of orange tentacles and heads of green ones in the A -regions, but they will lie at the tails of green and heads of orange in the B -regions. Examples of portions of these squares are shown in the figures in this section. See, for example, Figure 17.

To prove Theorem 4.1, we start with an EPD and pull it into a normal square following Lemma 4.2, and then analyze such normal squares. During the course of the proof, we will see that the first 2-edge loop near a white face of the square has one of seven possible forms, shown in Figures 10 through 16. These are analogous to Types \mathcal{A} through \mathcal{G} in [3, Figure 6.1], although more complicated behavior occurs here. In that paper, all regions were A -regions; in this one, regions may switch from A -regions to B -regions and vice versa. We record the results in a proposition.

Proposition 4.4. *Suppose E is an EPD in a prime diagram of a link K with an adequate, homogeneous state σ . Then the corresponding normal square E' of Lemma 4.2 has boundary running over tentacles adjacent to a 2-edge loop. Moreover, it has a white side whose nearest 2-edge loop has one of the forms in Figures 10 through 16, where in each general case, the dots indicate a staircase (with possibly no segments) of one of the forms of Lemma 4.10, shown in Figure 18.*




FIGURE 11. Type \mathcal{B} : EPD runs over tentacles of same color on 2-edge loop, white side cuts of triangle with vertex (zig-zag) containing one segment in the loop.




FIGURE 12. Type \mathcal{C} : Tentacles of different color, white side cuts off vertex not containing a segment of 2-edge loop, not separated by non-prime arc.




FIGURE 13. Type \mathcal{D} : Tentacles of same color, white side cuts off vertex not containing segment of 2-edge loop, and not separated from 2-edge loop by non-prime arc.

First, we deal with the case that the two white faces of the normal square both lie in the same polyhedral region. This case is simpler, because we may map into the corresponding lower polyhedron using a map called the “clockwise” or “counter-clockwise” map, and then apply standard results on alternating diagrams. We obtain the following.




FIGURE 14. Type \mathcal{E} : Tentacles opposite color, white side separated from 2-edge loop by non-prime arc.




FIGURE 15. Type \mathcal{F} : Tentacles of same color, white side separated from 2-edge loop by non-prime arc on same side of state circle as vertex (zig-zag).




FIGURE 16. Type \mathcal{G} : Tentacles of same color, white side separated from 2-edge loop by non-prime arc, non-prime arc on opposite side of state circle as vertex (zig-zag).

Lemma 4.5 (White faces in same region). *Suppose $D(K)$ is a prime diagram with an adequate, homogeneous state σ and corresponding upper polyhedron P_σ . Suppose E is an EPD in P_σ , and E' is the corresponding normal square from Lemma 4.2, intersecting white faces V and W in arcs β_V and β_W , respectively. If V and W are in the same polyhedral region, then $\partial E'$ runs over tentacles of a 2-edge loop in the upper polyhedron of type \mathcal{B} , Figure 11.*

The proof is nearly identical to that of [3, Lemma 6.7]. We step through it briefly to pick up the form of the 2-edge loop.

Proof. Observe that since V and W are in the same polyhedral region and our diagram is σ -homogeneous, then the region will either have all A or all B resolutions.

Apply the clockwise map of [3, Lemma 4.8] for an A -region, or the counter-clockwise map for a B -region (see also [9, Lemma 3.2] for both). By the cited results, we obtain a square S' in the lower polyhedron cutting off a single vertex in each white side. This square is either inessential, meaning it bounds a single ideal vertex in the lower polyhedron, or it can be isotoped to an EPD in the lower polyhedron. In the case that the square is essential, [3, Lemma 5.1] implies that S' runs over two segments of H_σ corresponding to a 2-edge loop. In the case that S' is inessential, its white arcs must cut off the same ideal vertex with opposite orientation, which implies one of the arcs lies on a bigon face. In either case, S' runs over segments of H_σ corresponding to a 2-edge loop. (See, e.g. [3, Figure 6.3].)

When the region is an A region, the proof of [3, Lemma 6.7] shows that E' , the original square, runs adjacent to a 2-edge loop as in the left of Figure 6.3 of that paper. When the region is a B region, the argument still applies, although the figure must be reflected and colors interchanged. In either case, $\partial E'$ runs adjacent to two tentacles of the same color (orange for A , green for B), running upstream. This is type \mathcal{B} , elementary, shown on the top left of Figure 11. \square

The next two lemmas help us identify a 2-edge loop when the EPD runs adjacent to distinct points on a state circle. We will apply them several times.

The first is identical to [3, Lemma 6.8], both in its statement and its proof, using the fact that the Utility Lemma, Lemma 3.5, and the Shortcut Lemma, Lemma 3.3, both hold in the homogeneously adequate setting.

Lemma 4.6 (Adjacent loop). *Let ζ be a directed simple arc contained in a single shaded face, adjacent to a state circle C at a point p on C . Suppose ζ runs upstream across a state circle C' after leaving p , but then eventually continues on to be adjacent to C again at a new point p' . Then ζ must run adjacent to two distinct segments of H_σ connecting C to C' .* \square

The next lemma is identical to [3, Lemma 6.9], also in its statement and proof. We include a proof here that expands a bit on the first sentence of the proof there, for additional clarity.

Lemma 4.7 (Adjacent points). *Suppose there are arcs ζ_1 and ζ_2 in distinct shaded faces in the upper polyhedron such that each ζ_j runs adjacent to two points p_1 and p_2 on the same state circle C (i.e. ζ_1 runs through a neighborhood of p_1 and a neighborhood of p_2 , and similarly for ζ_2). Then either*

- (1) *at least one of ζ_1 or ζ_2 runs upstream across some other state circle and Lemma 4.6 applies, or*
- (2) *both arcs remain adjacent to the same portion of C between p_1 and p_2 .*




FIGURE 17. Possibilities for Lemma 4.8. Two are in A -regions, two in B -regions. We refer to the forms as *hanging regions* and *standing regions*.

Proof. If one of ζ_1, ζ_2 crosses another state circle C' between points p_1 and p_2 then it must do so running upstream, by the Utility Lemma, Lemma 3.5. That is, if ζ_2 crosses C' after leaving p_1 , then it must cross C' again to reach p_2 . The Utility Lemma implies the first crossing is in the upstream direction. Then we are in case (1), and Lemma 4.6 applies.

So suppose neither ζ_1 nor ζ_2 crosses another state circle between points p_1 and p_2 . Then we form a simple closed curve meeting H_σ exactly twice at p_1 and p_2 by connecting portions of ζ_1 and ζ_2 between them. Replacing segments of H_σ with crossings, this gives a simple closed curve in the diagram of the link meeting the link transversely in exactly two points. Since the diagram is prime, as in Definition 2.7, the curve bounds no crossings on one side. It follows that ζ_1 and ζ_2 must run adjacent to C on this side. \square

We now return to proving Proposition 4.4. When we push an EPD into a normal square as in Lemma 4.2, the boundary of the square has two white sides, β_V and β_W , and two shaded sides. By Lemma 4.5, the proposition will be true if β_V and β_W lie in the same polyhedral region. We need to show the proposition is true if the white sides are in different polyhedral regions. In that case, the following lemma, the analogue of [3, Proposition 6.10], shows that β_V and β_W have a particular form.

Lemma 4.8 (Hanging and standing regions). *With the hypotheses of Proposition 4.4, either*

- (1) *the EPD runs adjacent to segments of a 2-edge loop of one of the desired types, or*
- (2) *the polyhedral regions containing β_V and β_W are of one of four forms, shown in Figure 17. The two forms on the left of that figure are referred to as hanging regions, the two on the right as standing regions, and we use A or B to denote whether the region is an A or B region.*

The proof of Lemma 4.8 follows that of [3, Proposition 6.10], which is easily generalized to the homogeneously adequate case. That proof is divided into four cases, but in all cases, everything of interest happens within the same polyhedral region. Because our diagram is homogeneous, all resolutions in the region must

be either all- A or all- B . If all- A , the proof of [3, Proposition 6.10] goes through without change. If all- B , the result holds by reflection. We walk through the argument briefly in order to determine the forms of the squares that arise.

Proof. First, by Lemma 4.5, we may assume white sides of $\partial E'$, β_V and β_W , are in distinct polyhedral regions. Then apply [3, Lemma 4.16]. This lemma is stated in the all- A case in [3], but it concerns one fixed polyhedral region. Thus it holds without change if our region is an A -region, and holds by reflection if our region is a B -region (this is noted in [3, Section 4.5]). The lemma states that when we direct arcs in the shaded faces toward β_W , one of them, call it ζ_1 , runs downstream across a state circle C_W , and then connects to β_W without crossing any additional state circles or non-prime arcs. The other, call it ζ_2 , may either connect directly to β_W when it enters the region of β_W , or it may run upstream first, across a different state circle C_2 , then back downstream later across C_2 to connect to β_W .

Two of the cases now concern ζ_2 , whether it connects directly to β_W or runs upstream first. In each case, β_W meets the tentacles containing ζ_1 and ζ_2 . We need to consider whether the head of the tentacle of ζ_1 meets the tail of that of ζ_2 , or whether the tail of the tentacle of ζ_1 meets the head of that of ζ_2 . Thus there are four cases total. We step through the cases briefly. Refer to [3] for complete details.

Case 1a: The arc ζ_2 runs directly to β_W and the head of the tentacle of ζ_2 meets the tail of the tentacle containing ζ_1 . Then ζ_1 and ζ_2 both run adjacent to segments connecting some C' to C_W , and these segments must be distinct by the fact that the diagram is prime. We pick up a 2-edge loop in this case that has the form of type \mathcal{A} , elementary, Figure 10.

Case 1b: The arc ζ_2 runs directly to β_W and the tail of the tentacle of ζ_2 meets the head of the tentacle of ζ_1 . This leads to one of the two standing pictures, depending on whether the region is an A -region or a B -region.

Case 2a: The arc ζ_2 runs upstream before running to β_W , and the head of the tentacle of ζ_1 meets the tail of the tentacle of ζ_2 . In this case, ζ_2 runs upstream adjacent to a segment joining C_W and some state circle C_2 . When ζ_2 runs downstream again across C_2 , it must run adjacent to another segment joining C_2 and C_W . These two segments form the desired 2-edge loop. It has the form of type \mathcal{B} , elementary, in Figure 11.

Case 2b: The arc ζ_2 runs upstream before running to β_W , and the tail of the tentacle of ζ_1 meets the head of the tentacle of ζ_2 . Again when ζ_2 runs upstream, it must do so adjacent to a segment s_2 connecting C_W and some C_2 . After ζ_2 runs downstream across C_2 , it connects immediately to β_W , implying that the tail of the tentacle of ζ_1 lies on C_2 . Then ζ_1 must run downstream adjacent to a segment s_1 connecting C_W and C_2 . If s_1 and s_2 are distinct, we have a 2-edge loop of the form of type \mathcal{A} .

If s_1 and s_2 agree, then notice that ζ_1 and ζ_2 both run adjacent to a point p on the state circle C_2 where $s_1 = s_2$ meets that state circle. They also both meet at β_W ,

hence by shrinking β_W , we may assume they run adjacent to the point p' on C_2 that is cut off by β_W . Apply Lemmas 4.6 (Adjacent Loop) and 4.7 (Adjacent points). If ζ_2 runs upstream from C_W , then Lemma 4.6 implies there is a 2-edge loop. Note it runs over tentacles of the same color. If ζ_2 runs from β_W in an A -region across C_2 to a B -region, or from β_W in a B -region across C_2 to an A -region, then the 2-edge loop will be of type \mathcal{B} , elementary, as on the 2nd row of Figure 11. If it runs from β_W in an A -region across C_2 to an A -region, or B -region to a B -region, it will be type \mathcal{D} , elementary, as in Figure 13, left.

If ζ_2 does not run upstream from C_W , Lemma 4.7 implies that the diagram has the form of one of the two hanging regions on the left of Figure 17. \square

The previous results imply that if there is not a 2-edge loop, β_V and β_W lie in different polyhedral regions and the regions in which they lie have one of the forms of Figure 17. We set up some notation. For each of the four forms of Figure 17, there is a horizontal state circle at the bottom of the diagram. In the case we are looking at the region of β_W , call that state circle C_W . If we are looking at β_V , call that state circle C_V . Note that C_W might equal C_V , but we can assume that the regions of β_W and β_V are distinct by Lemma 4.5.

In the following lemmas, we use the forms of Figure 17 to try to determine the form of adjacent regions, and eventually show there must be a 2-edge loop. The next two lemmas are the generalization of [3, Lemmas 6.11, 6.12] to the homogeneously adequate case.

Lemma 4.9. *Suppose $C_W \neq C_V$ and that ζ_2 , directed away from β_W , crosses C_W running downstream. Then for each ζ_i , $i = 1, 2$, there exists a staircase whose top is on C_W , with at least one segment, and with the following additional properties:*

- (1) *The arc ζ_i runs adjacent to each segment of the staircase it defines.*
- (2) *Both staircases run between C_W and a state circle C , where either C bounds a non-prime half-disk D if D is the first non-prime half-disk that ζ_2 runs into without exiting, or $C = C_V$ if ζ_2 exits every non-prime half-disk it enters.*
- (3) *Although the staircases of ζ_1 and ζ_2 may have distinct segments, the segments connect the same pairs of state circles each step of the way.*

Proof. To obtain the staircase of ζ_2 , apply Lemma 3.4 (Downstream) to ζ_2 directed away from β_W . Since ζ_2 runs downstream across C_W , there is at least one segment of the resulting staircase, ζ_2 runs adjacent to each segment, and the staircase either reaches C_V , or ends on the first state circle C for which ζ_2 crosses into a non-prime half-disk adjacent to C and does not exit.

Now consider the arc ζ_1 . If ζ_2 crosses into a non-prime half-disk without exiting, then the state circle C bounding that half-disk does not separate β_V and β_W . Hence note in this case that ζ_1 must also cross into the non-prime half-disk, and it must do so by crossing C . Since C is not separating, ζ_1 must actually cross it twice, once upstream and once downstream, by Lemma 3.5 (Utility). If, however, the staircase

of ζ_2 runs to β_V without running into a non-prime half-disk separating β_V and β_W , then ζ_2 crosses C_V running downstream. It follows from [3, Lemma 4.16] that ζ_1 must cross C_V running upstream. In all cases, ζ_1 , directed from β_W to β_V , crosses the last state circle in the staircase of ζ_2 (hereafter referred to as C) in the upstream direction.

To obtain the staircase of ζ_1 , change the direction on ζ_1 to run downstream from C to C_W and apply Lemma 3.4 (Downstream). This finishes the first two claims of the lemma.

To prove the final claim, the staircases of ζ_1 , ζ_2 , and arcs between their tops on C_W and on C define a closed curve γ in the graph H_σ . This curve crosses a state circle between every pair of segments, hence γ must cross each such state circle twice, and so the state circles connect in pairs within the region bounded by γ . By σ -adequacy of the diagram, no state circle can connect to the next adjacent state circle in the staircase. This forces the two staircases to have the same number of stairs, and the state circles to connect directly left to right, giving the third claim of the lemma. \square

Lemma 4.10 (Full staircase). *Suppose $C_W \neq C_V$, and that ζ_2 crosses C_W running downstream when directed from β_W to β_V . Then either:*

- (1) *the EPD runs over a 2-edge loop of one of the forms of Proposition 4.4; or*
- (2) *β_W is one of four forms:*
 - (a) *as on the far left of Figure 17 (hanging-A), with a nonempty sequence of segments forming a right-down staircase below C_W (i.e. all segments in A-regions), and with ζ_1 and ζ_2 adjacent to either side of each segment and each state circle of that staircase,*
 - (b) *as on second to left of Figure 17 (hanging-B) with a nonempty sequence of segments forming a left-down staircase below C_W (i.e. all segments in B-regions), and with ζ_1 and ζ_2 adjacent to either side of each segment and each state circle of the staircase,*
 - (c) *as on the second to right of Figure 17 (standing-A), with a nonempty left-down staircase (B-staircase), and with ζ_1 and ζ_2 adjacent to either side,*
 - (d) *as on the far right in Figure 17 (standing-B) with a nonempty right-down staircase (A-staircase) with ζ_1 and ζ_2 adjacent on either side.*

In the second case, the staircase is maximal in the sense that it either runs to C_V , or runs to the first state circle C for which there is a non-prime half-disk on C separating β_W and β_V .

The options are shown in Figure 18.

Proof. By Lemma 4.9, there exists a staircase from C_W either to C or C_V . Let k be the number of stairs traversed from C_W in the staircase. We will show by induction




FIGURE 18. Forms of staircases from Lemma 4.10. We say the right-down staircases on either end are A -staircases, and the left-down staircases in the center are B -staircases.

on k that either there is a 2-edge loop of one of the proper forms, or the $(k + 1)$ -st stair has the form of case (2) of the statement of the lemma.

For the base case of the induction, note that if we have traversed zero stairs of the staircase, that β_W must have one of the four forms of Figure 17 by Lemma 4.8. We may think of these as having one of the forms in Figure 18 when there are zero steps. (Note Lemma 4.9 will imply $k > 0$.)

So suppose after traversing $k \geq 0$ stairs there is no 2-edge loop, and β_W has one of the four forms in the statement of the lemma, possibly with an empty staircase if $k = 0$. Let s_1 denote the $(k + 1)$ -st segment of the staircase of ζ_1 , and let s_2 denote the $(k + 1)$ -st segment of the staircase of ζ_2 . If $s_1 \neq s_2$, then there is a 2-edge loop at this step of the staircases. Because the boundary of the EPD runs through tentacles of different color adjacent to the segments of the 2-edge loop, the 2-edge loop is of type \mathcal{A} or \mathcal{C} . If we have an A -staircase running to an A -region, or a B -staircase running to a B -region, it will be type \mathcal{A} , general. Otherwise, it will be type \mathcal{C} , general.

So suppose $s_1 = s_2$. Now, the bottom of the k -th segment meets C_k at a point p and ζ_1 and ζ_2 both run adjacent to p on C_k (or if $k = 0$ and β_W lies in a standing region, then p is the point at the end of the segment on C_W , and ζ_1 and ζ_2 both run adjacent to p). The top of the $(k + 1)$ -st segment meets C_k at a point p' , and ζ_1 and ζ_2 both run adjacent to p' . If one of ζ_1 , ζ_2 runs upstream, then Lemma 4.6 (Adjacent Loop) implies that this arc runs through a 2-edge loop. If ζ_2 runs through the 2-edge loop, then the loop is on the same side of C_k as the staircase from β_W , and it will be of type \mathcal{F} if $k > 0$ or β_W lies in a hanging region, and of type \mathcal{B} -elementary if $k = 0$ and β_W lies in a standing region. If ζ_1 runs through the 2-edge loop, then the 2-edge loop lies on the opposite side of C_k as the staircase from β_W , and it will be of type \mathcal{G} .

Suppose then that there is no 2-edge loop. Lemma 4.7 (Adjacent Points) implies that ζ_1 and ζ_2 run parallel to C_k between p and p' .

Now, if β_W is in a hanging- A region, by induction the k steps between C_W and C_k form a right-down staircase through A -regions (case (2)(2a)). Arcs ζ_1 and ζ_2 run parallel to C_k from the end of the k -th segment to the beginning of the $(k + 1)$ -st.

The arc ζ_2 , in the green face, is running to the right across the top of C_k . If the $(k+1)$ -st segment of the staircase is in a B -region, then the arc ζ_1 , in the orange face, must run to the left across the bottom of C_k . This contradicts the fact that the arcs run parallel. So the $(k+1)$ -st segment is in an A -region, and ζ_1, ζ_2 run parallel to each other from the base of the k -th step to the base of the $(k+1)$ -st step. Thus the staircase is still of the form in case (2)(2a). The same argument, up to reflection, applies if β_W lies in a hanging- B region.

Now suppose that β_W lies in a standing- A region, and the staircase from C_W to C_k is of the form in case (2)(2c). That is, each segment runs through a B region. Again if there is no 2-edge loop, then ζ_1 and ζ_2 run parallel to C_k from the end of the k -th segment to the beginning of the $(k+1)$ -st segment. The arc ζ_2 , in the orange face, runs from the bottom of the k -th segment to the left across the top of C_k . Thus the arc ζ_1 , in the green face, must run to the left across the bottom of C_k . This will happen only if the region containing the $(k+1)$ -st segment is a B -region, giving the inductive step in this case. The same argument, up to reflection, gives the inductive step in the case β_W lies in a standing- B region. \square

Lemma 4.11 (Terminating staircase). *Suppose ζ_2 runs across C_W in the downstream direction when directed from β_W to β_V , out of every non-prime half-disk that it enters, and terminates with ζ_2 crossing C_V . Then the conclusion of Proposition 4.4 holds.*

Proof. By Lemma 4.10, we may assume that either the result holds, or $C_V = C_W$, or β_W has one of the four forms of Lemma 4.10 (Staircase). Swapping the roles of V and W , in the latter case the lemma implies that β_V also has one of those four forms.

Suppose first that $C_V \neq C_W$ and that β_W lies in a hanging- A region. Then there is a nonempty right-down staircase from β_W to β_V with ζ_1 and ζ_2 adjacent to each segment and each state circle; segments run only through A -regions, as in (2a) of Figure 18. Similarly, Lemma 4.10 implies there is such a staircase from β_V to β_W . This must agree with the former staircase, so each segment in it must run only through A -regions. So β_V is either in a hanging- A or standing- B region, and β_V is attached to a right-down staircase. However, note that in these two cases, the green (darker shaded) face will be adjacent to the side of C_V that contains the staircase from β_W , since green runs along “tops” of state circles, i.e. the side containing β_W , on the right-down staircase from β_W in a hanging- A region, item (2a) of Figure 18. The green face will also be adjacent to the opposite side of C_V , since if β_V lies in a hanging- A or standing- B region, the green face is adjacent to the side of the state circle C_V containing β_V . But then we may draw an arc in the green face with one endpoint on one side of C_V , and one endpoint on the other, which is disjoint from C_V in the interior. This arc can be connected in a neighborhood of C_V to a closed

curve in H_σ meeting C_V exactly once, which is a contradiction. The same argument, reflected, implies that β_W does not lie in a hanging- B region if $C_W \neq C_V$.

So now suppose $C_V \neq C_W$ and β_W lies in a standing- A region. Then ζ_1 and ζ_2 run adjacent to a nontrivial left-down staircase through B -regions. Again, β_V must also give a staircase through B regions, so β_V either lies in a hanging- B or standing- A region. Again in either case we have orange faces adjacent to either side of C_V , giving a contradiction as above. A reflection of this argument implies that β_W does not lie in a standing- B region if $C_V \neq C_W$.

Thus $C_V = C_W$. We still know that β_V and β_W are either in hanging or standing regions, one on either side of $C := C_W = C_V$. To avoid having green or orange faces adjacent to both sides of C , we must have one of the combinations: hanging- A with standing- A , hanging- A with hanging- B , standing- A with standing- B , and hanging- B with standing- B .

In all cases, consider the point on C where the segment of a hanging or standing region meets C . In the hanging case, both ζ_1 and ζ_2 run adjacent to this point. In the standing case, slide the point slightly to lie between ζ_1 and ζ_2 on C , and again we have ζ_1 and ζ_2 running adjacent to this point. Then in all possible combinations, Lemma 4.7 (Adjacent Points) implies that either there is a 2-edge loop, or ζ_1 and ζ_2 run parallel to C between these points. If there is a 2-edge on the same side of C as a hanging region, then it must be of type \mathcal{F} . If there is a 2-edge loop on the same side of C as a standing region, it must be of type \mathcal{B} -elementary.

So suppose that ζ_1 and ζ_2 run parallel to C between the two points. Consider the hanging- A with standing- A case, sketched with C horizontal, with the hanging region on top. Note that at the base of the segment on C in the hanging region, the arc in the green face runs to the right adjacent to C . However, on the opposite side, the arc in the orange face runs to the left. This is a contradiction. We obtain the same contradiction for hanging- A with hanging- B , standing- A with standing- B , and for hanging- B with standing- B . So in all cases, there must be a 2-edge loop. \square

Lemma 4.12 (Inside non-prime arcs). *Suppose β_W and β_V are separated by a non-prime arc α , with the arc ζ_2 , say, crossing α . Suppose α is outermost among all such arcs, with respect to β_W . That is, α is the first such non-prime arc crossed by ζ_2 when directed toward β_V . Then we have the conclusion of Proposition 4.4.*

Proof. As in the proof of [3, Lemma 6.14], we break into two cases: whether ζ_2 runs upstream after crossing α or not.

Case 1. Suppose ζ_2 does not run upstream after crossing α . Let C be the state circle on which the endpoints of α lie. First, note that ζ_2 cannot run downstream across C after crossing α , for since β_V lies inside the non-prime half-disk bounded by α and C , if ζ_2 crosses C running downstream away from β_V , it must cross C again to reach β_V . This contradicts the Utility Lemma 3.5.

Since ζ_2 does not run upstream or downstream after crossing α , it must run to the vertex β_V . We know that β_V has one of the four forms in Figure 17. Since ζ_2 does not run upstream or downstream, but both arcs in the hanging regions run upstream or downstream, β_V must be as on the right of that figure, one of the two standing regions. We will assume it is a standing- A region, since the standing- B argument is the same up to reflection. Thus ζ_2 lies in an orange face.

Lemma 4.10 (Full staircase) implies that on the opposite side of α , either there is a 2-edge loop, or ζ_1 and ζ_2 run adjacent to a (possibly empty) staircase of β_W . If the staircase of β_W is non-empty, then it has one of the four forms given by that lemma. To ensure the colors of faces match correctly, ζ_2 must run through an orange face, and so the staircases must either be type (2c), standing- A and left-down, or type (2b), hanging- B and left-down. But now note that a left-down staircase terminates in a B -region, at the state circle C , whereas we are assuming α has endpoints in an A -region, on the same side of the state circle C . This contradicts our assumption of σ -homogeneity.

Thus we reduce to the case that the staircase of β_W is empty, or consists of no segments. The same argument as above implies that β_W does not lie in a hanging region. Thus β_W also lies in a standing region, and to ensure homogeneity, it must be a standing- A region. We assume ζ_2 meets no state circles when running from β_W to β_V , else we'll get a 2-edge loop (of type \mathcal{B} -elementary) by Lemma 4.6 (Adjacent loop). Additionally, ζ_1 crosses C twice, and we may assume it meets no state circles between crossings, or again Lemma 4.6 (Adjacent loop) implies that there is a 2-edge loop, this time of type \mathcal{D} -elementary or \mathcal{B} -elementary. Then $\zeta_1 \cup \zeta_2$ gives a closed curve in the graph H_σ meeting H_σ exactly twice. This gives a closed curve meeting the diagram twice with crossings on either side (inside and outside the half-disk bounded by α), contradicting the fact that the diagram is prime, Definition 2.7.

Case 2. The arc ζ_2 runs upstream after crossing α .

If ζ_2 runs back to C afterward, then there must be a 2-edge loop by Lemma 4.6 (Adjacent loop), of type \mathcal{F} in Figure 15. So suppose ζ_2 does not return to C after running upstream, across a state circle we will label C' .

We claim that ζ_1 must also cross C' . If C' separates β_V and β_W , then both ζ_1 and ζ_2 must cross C' . So suppose C' does not separate β_V and β_W . Then ζ_2 must cross C' twice, the second time in the downstream direction, running toward β_V . By assumption, ζ_2 does not run back to C , hence it runs downstream to some C'' . The arc ζ_1 crosses C running downstream toward β_V , so it must also run downstream to C'' , following tentacles to the same side of C'' as ζ_2 . But this leads to a contradiction: such arcs cannot meet in a vertex, but will be separated by segments. So we conclude that both ζ_1 and ζ_2 run along segments between C and C' . If these segments are distinct, then there is a 2-edge loop, of type \mathcal{E} .

If these segments are not distinct, then both ζ_1 and ζ_2 run adjacent to the endpoint p of the segment on C . They also both run adjacent to a point p' on C where an

endpoint of α meets C . Between those two points, ζ_2 runs over α , hence meets no segments or state circles. If the arc ζ_1 meets state circles between, then Lemma 4.6 (Adjacent loop) implies there is a 2-edge loop. This will either be of type \mathcal{B} -general, of type \mathcal{D} -general, or of type \mathcal{G} , depending on whether ζ_1 runs into an A or B -region, and whether it runs over a non-prime arc.

If ζ_1 meets no state circles in between p and p' , then take the portions of ζ_1 and ζ_2 between these two points and connect their ends. This gives a loop meeting the diagram exactly twice, with crossings on either side, contradicting the fact that the diagram is prime. \square

Proof of Proposition 4.4. Let E be an essential product disk embedded in P_σ . By Lemma 4.2, there exists a normal square E' with each white side cutting off a triangle in such a way that in counter-clockwise order, edges of the triangle are colored white–orange–green. As usual, label the white sides β_V and β_W . By Lemma 4.5, if β_V and β_W are in the same polyhedral region, then there is a 2-edge loop in the upper polyhedron with one of the desired forms. If β_V and β_W are not in the same region, then the sides of $\partial E'$ in the shaded faces, which we label ζ_1 and ζ_2 , must run over a sequence of state circles to get from β_W to β_V . Lemma 4.11 implies that if ζ_2 crosses C_W and runs out of every non-prime half disk until it crosses C_V , then the boundary of E' runs over tentacles adjacent to segments of a 2-edge loop as desired. If instead ζ_2 runs into a non-prime half-disk without exiting, then Lemma 4.12 implies that the boundary of E' runs over tentacles adjacent to a 2-edge loop as desired. \square

The proof of Theorem 4.1 is now immediate.

Proof of Theorem 4.1. Given an EPD E , pull it into a normal square E' as in Lemma 4.2. By Proposition 4.4, the normal square has one of the forms of Figures 10 through 16. In every case, the boundary of the square E' runs through tentacles adjacent to segments of a 2-edge loop. We may slide the white sides β_V and β_W of $\partial E'$ back onto the ideal vertices to recover an EPD isotopic to E , and we may do this in such a way that ∂E runs over the same 2-edge loop as $\partial E'$, concluding the proof of the theorem. \square

5. VOLUME APPLICATIONS

In this section, we apply of Theorem 4.1 to bound volumes of classes of links in terms of their diagrams. First, we need some definitions.

Definition 5.1. A *twist region* is a string of bigon regions in the complement of the diagram graph that are adjacent end to end, or a single crossing adjacent to no bigons. We also require the diagram to be alternating in a twist region, else a Reidemeister move simplifies the diagram, and we require that a twist region be maximal in the sense that there are no additional bigons adjacent on either side. An example is shown in Figure 19, left.




FIGURE 19. Left to right: A twist region, its short resolution, and its long resolution.

Given a twist region with more than one crossing, an adequate state will choose either all A or all B resolutions for that twist region. One choice will produce a portion of H_σ consisting of two state circles, with segments parallel between them. The other choice will produce a sequence of state circles from the interior of the twist region, each meeting only two segments. We call the former the *short* resolution, and the latter the *long* resolution. See Figure 19.

The following definition is as in [4].

Definition 5.2. Two edges of a 2-edge loop in \mathbb{G}_σ belong to the same twist region R of the diagram if the edges correspond to crossings of R , and the resolution of R is the short one. We say that \mathbb{G}_σ satisfies the *2-edge loop condition* if every 2-edge loop in \mathbb{G}_σ belongs to the same twist region. In this case, we also say that the diagram satisfies the 2-edge loop condition for σ .

Theorem 4.1 gives us the following corollary, analogous to [3, Corollary 5.19].

Corollary 5.3. Suppose that $D(K)$ is a prime diagram with homogeneous, adequate state σ , and suppose that \mathbb{G}_σ satisfies the 2-edge loop condition. Then $\|E_c\| = 0$.

Proof. Recall from Definition 2.19 that $\|E_c\|$ is the number of complex EPDs in a minimal spanning set for the I -bundle of $S^3 \setminus S_\sigma$, from [3, Lemma 5.8]. Let R be a twist region of $D := D(K)$ whose resolution under σ is the short resolution. In [3, Lemma 5.17], which extends immediately to the homogeneously adequate case (see Section 5.6 of that paper), it is shown that if \hat{D} is the diagram obtained by removing one crossing from R , then $E_c(D)$ and $E_c(\hat{D})$ have the same cardinality.

Now let \bar{D} be the diagram obtained from D by removing all bigons in the short resolutions of twist regions of D . By induction, $\|E_c(D)\| = \|E_c(\bar{D})\|$. Moreover, because \mathbb{G}_σ satisfies the 2-edge loop condition, this will remove all 2-edge loops from $\mathbb{G}_\sigma(\bar{D})$.

But by Theorem 4.1, any EPD coming from \bar{D} has boundary running over a 2-edge loop. Since $\mathbb{G}_\sigma(\bar{D})$ has no 2-edge loops, there must be no EPDs in $E_c(\bar{D})$. Hence $\|E_c\| = 0$. \square

Suppose that M is a 3-manifold and S is an essential surface embedded in M . Recall that the *guts* of $M \setminus S$ is defined to be the complement of the maximal I -bundle in $M \setminus S$. Work of Agol [1], extended by Kuessner [10], says that guts can be used to bound the Gromov norm of M . In particular,

$$(1) \quad \|M\| \geq 2 \chi_-(\text{guts}(M \setminus S)).$$

This was extended by Agol, Storm, and Thurston to give bounds on the volumes of hyperbolic manifolds [2]. If M is hyperbolic, then

$$(2) \quad \text{vol}(M) \geq v_8 \chi_-(\text{guts}(M \setminus S)),$$

where $v_8 = 3.6638 \dots$ is the volume of a regular ideal octahedron.

In [3], it was shown that

$$\chi_-(\text{guts}(S^3 \setminus S_\sigma)) = \chi_-(\mathbb{G}'_\sigma) - \|E_c\|,$$

where recall \mathbb{G}'_σ is the reduced state graph of Definition 2.2. This is Theorem 5.14 of that paper, extended to homogeneously adequate links in [3, Section 5.6]. Combining this with Corollary 5.3 above, we obtain the following corollary, analogous to [3, Corollary 9.4].

Corollary 5.4. *Suppose that $D(K)$ is a prime diagram of a hyperbolic link K , with homogeneous, adequate state σ , and suppose that \mathbb{G}_σ satisfies the 2-edge loop condition. Then*

$$\text{vol}(S^3 - K) \geq v_8 (\chi_-(\mathbb{G}'_\sigma)).$$

Proof. By Theorem 5.14 of [3], which applies to the homogeneously adequate case, $\chi_-(\text{guts}(S^3 \setminus S_\sigma)) = \chi_-(\mathbb{G}'_\sigma) - \|E_c\|$ where $\|E_c\|$ is the number of complex essential product disks in a spanning set for the I -bundle. By Corollary 5.3 above, $\|E_c\| = 0$. Now equation (2) above implies that the volume of $S^3 - K$ is bounded below by $v_8 \chi_-(\mathbb{G}'_\sigma)$. \square

REFERENCES

1. Ian Agol, *Lower bounds on volumes of hyperbolic Haken 3-manifolds*.
2. Ian Agol, Peter A. Storm, and William P. Thurston, *Lower bounds on volumes of hyperbolic Haken 3-manifolds*, J. Amer. Math. Soc. **20** (2007), no. 4, 1053–1077, with an appendix by Nathan Dunfield.
3. David Futer, Efstratia Kalfagianni, and Jessica Purcell, *Guts of surfaces and the colored Jones polynomial*, Lecture Notes in Mathematics, vol. 2069, Springer, Heidelberg, 2013. MR 3024600
4. David Futer, Efstratia Kalfagianni, and Jessica S. Purcell, *Hyperbolic semi-adequate links*.
5. ———, *Dehn filling, volume, and the Jones polynomial*, J. Differential Geom. **78** (2008), no. 3, 429–464.
6. ———, *Symmetric links and Conway sums: volume and Jones polynomial*, Math. Res. Lett. **16** (2009), no. 2, 233–253.
7. ———, *On diagrammatic bounds of knot volumes and spectral invariants*, Geom. Dedicata **147** (2010), 115–130. MR 2660569
8. ———, *Jones polynomials, volume, and essential knot surfaces: a survey*, Proceedings of Knots in Poland III, vol. 100, Banach Center Publications, 2014, pp. 51–77.
9. ———, *Quasifuchsian state surfaces*, Trans. Amer. Math. Soc. **366** (2014), no. 8, 4323–4343. MR 3206461
10. Thilo Kuessner, *Guts of surfaces in punctured-torus bundles*, Arch. Math. (Basel) **86** (2006), no. 2, 176–184. MR 2205233 (2006k:57049)

11. Marc Lackenby, *The volume of hyperbolic alternating link complements*, Proc. London Math. Soc. (3) **88** (2004), no. 1, 204–224, With an appendix by Ian Agol and Dylan Thurston. MR MR2018964 (2004i:57008)
12. W. B. Raymond Lickorish and Morwen B. Thistlethwaite, *Some links with nontrivial polynomials and their crossing-numbers*, Comment. Math. Helv. **63** (1988), no. 4, 527–539. MR MR966948 (90a:57010)
13. William W. Menasco, *Polyhedra representation of link complements*, Low-dimensional topology (San Francisco, Calif., 1981), Contemp. Math., vol. 20, Amer. Math. Soc., Providence, RI, 1983, pp. 305–325. MR MR718149 (85e:57006)
14. George D. Mostow, *Quasi-conformal mappings in n -space and the rigidity of hyperbolic space forms*, Inst. Hautes Études Sci. Publ. Math. (1968), no. 34, 53–104. MR 0236383 (38 #4679)
15. Makoto Ozawa, *Essential state surfaces for knots and links*, J. Aust. Math. Soc. **91** (2011), no. 3, 391–404.
16. Gopal Prasad, *Strong rigidity of \mathbf{Q} -rank 1 lattices*, Invent. Math. **21** (1973), 255–286.
17. Jessica S. Purcell, *Volumes of highly twisted knots and links*, Algebr. Geom. Topol. **7** (2007), 93–108. MR 2289805 (2007m:57011)
18. Morwen B. Thistlethwaite, *On the Kauffman polynomial of an adequate link*, Invent. Math. **93** (1988), no. 2, 285–296. MR 948102 (89g:57009)
19. William P. Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 3, 357–381.

MATHEMATICS DEPARTMENT, BRIGHAM YOUNG UNIVERSITY, PROVO, UT 84602, USA